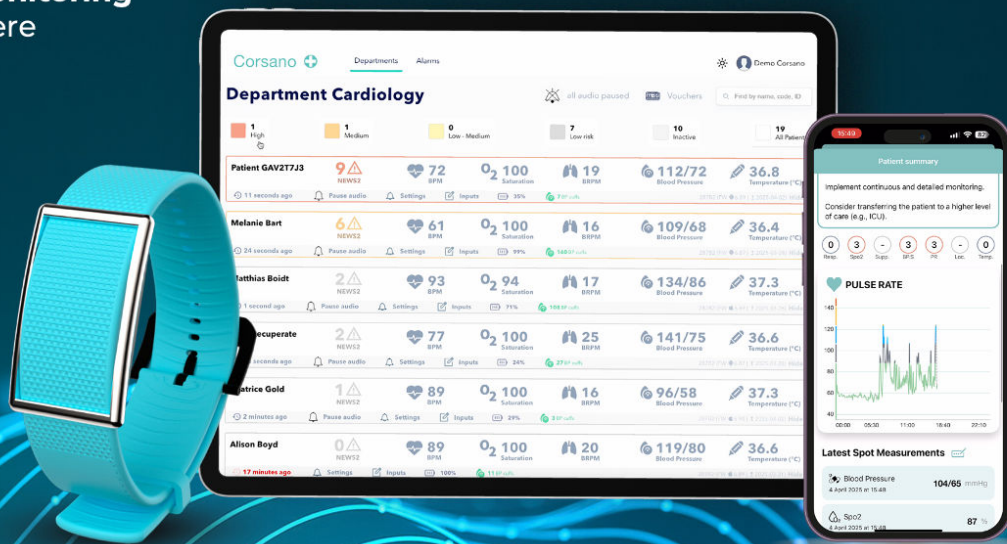


Continuous Patient Monitoring  
Anytime, Anywhere



## White Paper

### The Full-Stack Advantage

*Why Hardware Sensor Owners Are Best Positioned to Build a Meaningful Health OS*

A Critical Analysis of Raw Data Access, AI Training and Regulatory Compliance  
and the Strategic Value of Vertical Integration in Digital Health



## Executive Summary

The wearable health technology market is converging on a pivotal architectural question. Who is best positioned to build a Health Operating System (Health OS) capable of delivering clinically meaningful, AI-driven insights at scale? This paper argues that vertically integrated hardware sensor owners hold decisive structural advantages over agnostic platform aggregators. These companies control the full stack: from photoplethysmography (PPG) hardware and firmware through proprietary signal processing algorithms and cloud infrastructure, all the way to clinical-grade regulatory approvals.

The core hypothesis is straightforward. Proximity to raw physiological data is not merely a technical preference. It is a fundamental prerequisite for training AI models that go beyond what is achievable by recombining intermediate, pre-processed vital parameters. This paper also critically examines whether granular data is always necessary. For certain AI applications, aggregate data can suffice. Recognizing that distinction is essential for calibrating investment and development priorities.

Corsano Health B.V. occupies a uniquely strong position in this landscape. It is one of the few companies globally that combines: (1) a clinically-validated, MDR Class IIa and FDA 510(k)-cleared continuous monitoring wearable producing data at up to 128 Hz; (2) proprietary firmware and algorithms for all major vital parameters; (3) a full software stack including patient and HCP portals; and (4) strategic partnerships with Medtronic, Netcare and major academic hospitals. Corsano is a rare case of full-stack ownership in the medical-grade wearable segment. This paper explains why that architecture creates durable competitive advantages in the emerging Health OS paradigm.

### Core Thesis

Hardware sensor owners that control raw physiological signal data are structurally advantaged in building AI-powered Health OS platforms. However, the marginal value of granularity diminishes for certain population-level applications. A calibrated, tiered data architecture is required to capture both advantages.

## 1. Introduction: The Health OS Imperative

### 1.1 The Emerging Concept of a Health Operating System

The term 'Health OS' has gained traction as the conceptual equivalent of a personal operating system for physiological state. It is an always-on, continuously learning platform that ingests biological signals, contextual data and behavioral inputs to generate actionable health intelligence. Traditional electronic health records (EHRs) are episodic and retrospective. A Health OS is continuous, predictive and personalized. It is the metabolic and cardiovascular equivalent of a financial portfolio management system: real-time, data-rich and capable of trend recognition far ahead of symptomatic presentation.

Three categories of companies are competing to own this platform layer. The first is consumer wearable hardware companies such as Apple, Samsung, Garmin and Fitbit/Google. The second is specialist wearable companies with proprietary sensor platforms including Corsano, Oura, Whoop and Abbott. The third is platform aggregators such as Apple HealthKit, Google Health Connect and Samsung Health. These aggregators unify data streams from multiple hardware sources without owning the underlying signal generation.



The central competitive question is not who has the most users. It is who has access to the data architecture necessary to train AI models that produce meaningful, clinically valid insights. This paper analyzes that question through four lenses: signal quality and data granularity, AI training requirements, regulatory moats and full-stack control.

## 1.2 Why This Question Matters Now

The convergence of several structural trends makes this question urgent. Large language models (LLMs) and multimodal AI architectures have demonstrated that data quality and provenance are at least as important as model architecture in determining output quality. Regulatory agencies, particularly the FDA and EMA, are increasingly requiring clinical-grade evidence from AI-driven diagnostic and monitoring claims. This creates a bifurcation between wellness analytics and medical-grade insights.

The reimbursement landscape is also shifting. Payors and health systems are beginning to reward continuous monitoring with Remote Patient Monitoring (RPM) codes. This creates economic incentives that favor clinical-grade platforms over consumer wellness apps. Meanwhile, major technology companies including Apple, Google and Samsung have made large investments in health platform aggregation. That raises the question of whether specialist hardware companies can defend their data moat over the long term.

## 2. The Data Hierarchy: Raw Signals, Intermediate Parameters and Aggregate Insights

### 2.1 A Taxonomy of Wearable Health Data

Understanding the competitive dynamics of Health OS platforms requires a precise taxonomy of the data layers produced by wearable physiological monitoring systems. These layers are not equivalent. They carry fundamentally different information densities and the transformations between layers are lossy by design.

Layer	Description	Examples	Access
<b>L0 - Raw Signal</b>	Unprocessed optical or electrical time-series at full sensor resolution	PPG waveform at up to 128 Hz, raw ECG samples, accelerometer XYZ at 32 Hz	<b>Hardware owner only</b>
<b>L1 - Derived Parameter</b>	Computed vital sign from proprietary algorithm applied to L0	Heart rate, SpO2, HRV (RMSSD), respiration rate, skin temperature	Hardware owner + licensed SDK partners
<b>L2 - Composite Score</b>	Algorithm-generated wellness or health index from multiple L1 parameters	Oura Readiness Score, Whoop Recovery Score, Apple Activity Ring	Hardware owner; opaque to aggregators
<b>L3 - Aggregated Summary</b>	Temporally averaged or summarized L1/L2 data exposed via platform APIs	Daily step count, average resting heart rate via Apple HealthKit/Google Health	Agnostic platforms, third-party apps

Table 1. Wearable Health Data Hierarchy - Information Density and Access by Layer



## 2.2 Information Loss Across Layers: The Compression Problem

Each transition between data layers involves irreversible information compression. When a raw PPG waveform sampled at up to 128 Hz is processed to extract a single heart rate value, the majority of the physiological information content of the waveform is discarded. That discarded information includes pulse wave morphology, dicrotic notch characteristics, augmentation index, beat-to-beat variability at sub-second resolution and signal quality indicators. These are essential for distinguishing true physiological variation from motion artifact.

The compression is intentional and rational at the point of consumption. Users do not need access to 7,680 individual PPG data points per minute. But the compression creates a fundamental problem for AI training. The information that distinguishes a clinically significant arrhythmia from a clean sinus rhythm (or a genuine drop in oxygen saturation from a motion artifact) exists primarily at L0 and early L1. It does not exist at L2 or L3.

### Signal Fidelity Principle

At up to 128 Hz continuous sampling, the Corsano CardioWatch 287 generates approximately 11,059,200 PPG data points per day per patient. A typical HealthKit-aggregated daily summary contains fewer than 1,440 discrete heart rate readings: a data reduction ratio exceeding 99.98%. The discarded data contains the morphological signatures that AI models require to detect atrial fibrillation, early-stage hypoxemia and autonomic dysregulation.

## 2.3 The Intermediate Parameter Problem: Heterogeneity Across Manufacturers

A critical and underappreciated problem with platform aggregation approaches is not merely the data reduction from L0 to L3. It is the fundamental heterogeneity of the L1 parameters that different manufacturers compute from the same underlying physiological phenomenon. Heart rate variability (HRV) is the canonical example.

HRV can be expressed as RMSSD, SDNN, pNN50, LF/HF ratio or any number of frequency-domain metrics. Each manufacturer applies a different algorithmic pipeline to compute these values, using different peak detection algorithms, different artifact rejection strategies, different averaging windows and different normalization approaches. When an agnostic platform aggregates HRV values from a Whoop, an Oura Ring and a Garmin watch, it is combining three fundamentally incommensurable numbers that happen to share a label.

This heterogeneity is not a solvable data engineering problem. It is an inherent consequence of proprietary algorithmic differentiation. The only way to achieve algorithmically consistent vital parameter computation across a population is to control the full computation pipeline from raw signal to output. That requires owning the hardware, firmware and signal processing algorithms.

HRV Metric	Corsano	Oura	Whoop	Apple Watch
<b>Primary metric</b>	RMSSD (ms)	RMSSD (ms)	RMSSD (ms)	SDNN (ms)
<b>Measurement window</b>	Continuous / 30s epochs	Sleep-derived (3-5 min)	Light sleep, 5-min window	60-second breath focus
<b>Artifact rejection</b>	Proprietary multi-channel	Proprietary	Proprietary	Proprietary
<b>API exposure</b>	Full + raw PPG access	RMSSD only (HealthKit)	Score only (no HRV API)	SDNN via HealthKit
<b>Cross-device comparability</b>	N/A (single device)	<b>LOW (different window)</b>	<b>VERY LOW (opaque)</b>	<b>LOW (different metric)</b>

Table 2. HRV Measurement Heterogeneity Across Major Wearable Platforms



The implication for AI model training is profound. An LLM or predictive model trained on aggregated HRV data from HealthKit will learn from a composite signal that is algorithmically inconsistent. It effectively adds a hidden confounding variable (device-manufacturer identity) to every training observation. This creates models whose predictions are device-confounded rather than physiologically grounded.

## 3. Critical Analysis: Is Granular Data Always Better?

---

### 3.1 The Case for Granularity: Where Raw Data Creates Irreplaceable Value

The case for raw data access rests on specific use cases where granularity is not merely preferable but clinically necessary. These use cases define the ceiling of what a Health OS built on aggregated data can achieve.

#### 3.1.1 Arrhythmia Detection and Morphological Classification

Atrial fibrillation (AF) detection from a wrist-worn PPG sensor is the most well-validated wearable clinical use case to date. The algorithms that achieve clinically meaningful sensitivity and specificity (in the range required for FDA 510(k) clearance) operate on the inter-beat interval time series at full resolution combined with pulse wave morphology features extracted from the raw PPG waveform. Beat-to-beat RR interval variability at millisecond resolution is the signal of interest. A temporally averaged heart rate is not merely unhelpful: it actively obscures the signal. No amount of sophisticated model architecture can recover the AF signature from an hourly averaged heart rate.

Beyond binary AF detection, differentiating between AF, atrial flutter, premature ventricular contractions (PVCs) and sinus arrhythmia requires morphological features that exist exclusively at L0. For agnostic platforms, this is not a near-term capability gap. It is a structural impossibility.

#### 3.1.2 Early Sepsis and Deterioration Detection

The NEWS2 (National Early Warning Score 2) clinical framework requires continuous, high-resolution monitoring of six physiological parameters simultaneously: respiratory rate, oxygen saturation, systolic blood pressure, heart rate, consciousness level and temperature. Corsano has implemented this framework in its Cor-Events alert architecture. The clinical evidence for early deterioration detection is clear. The predictive window for sepsis identification extends to 12 to 24 hours before conventional clinical recognition. This advantage only holds when monitoring is continuous and high-frequency.

Intermittent sampling at the temporal resolution available from aggregated platform APIs (typically 1 to 5 minute averages) collapses this predictive window dramatically. Corsano's up to 128 Hz continuous sampling generates up to 7,680 discrete physiological measurements per minute across multiple channels. This provides the temporal resolution necessary to detect the subtle, progressive patterns (tachycardia trend plus respiratory rate elevation plus SpO2 drift) that constitute the early sepsis phenotype. This pattern is not detectable at 5-minute sampling intervals.



### 3.1.3 Sleep Architecture Analysis

Accurate sleep staging (distinguishing N1, N2, N3 and REM sleep) from a wrist-worn sensor requires access to continuous PPG morphology, accelerometry and ideally skin conductance or temperature signals at high frequency. Oura and Whoop have published internal validation studies for their sleep staging algorithms. However, these algorithms are proprietary. The derived sleep stage classifications that leave the device are fixed categorical labels (Light/Deep/REM) that cannot be reanalyzed with a different model.

A Health OS built on HealthKit sleep data is trained on the classification output of competitor algorithms. It cannot improve upon or customize the underlying sleep staging. Corsano's raw accelerometer and PPG data enables implementation of research-validated sleep staging algorithms directly. This includes the most current published approaches incorporating HRV-derived autonomic markers that post-date most competitor firmware. The result is a capability for algorithmic improvement that is structurally unavailable to agnostic platforms.

## 3.2 The Case Against Unconditional Granularity: Where Aggregate Data Is Sufficient

Intellectual honesty requires acknowledging that granular raw data does not provide universal advantage. There are important use cases where aggregate or intermediate data is adequate and where the cost and complexity of full raw data pipelines may not be justified.

### 3.2.1 Population-Level Epidemiological Modeling

For population health management applications such as identifying high-risk cohorts, predicting hospital readmission rates and stratifying patients by activity quartile, aggregate data is frequently sufficient. Public health studies using step count data from Fitbit at the population level ( $N > 100,000$ ) have demonstrated statistically robust relationships between daily activity and mortality risk. The granularity of a 128 Hz PPG signal adds no marginal value to a study asking whether people who walk more than 7,500 steps per day have lower 10-year cardiovascular mortality.

This is an important concession, but it comes with a critical caveat. Population-level models trained on aggregate data are not the same as individualized predictive models for a specific patient. The Health OS value proposition is fundamentally personalized. It requires individual-level prediction, not population-level association.

### 3.2.2 Behavioral and Lifestyle Coaching Applications

For corporate wellness applications such as the Corsano CardioMood platform targeting People & Culture professionals, the actionable insights are predominantly behavioral: sleep consistency, activity patterns and stress reactivity. These insights can be generated from L1 and L2 data with acceptable fidelity. A corporate wellness platform does not require beat-level PPG morphology to generate a recommendation to take a walking break at 2 PM. In this segment, the competitive advantage from raw data access is real but modest. The differentiation shifts toward algorithm quality, personalization depth and user experience.

### 3.2.3 Longitudinal Trend Monitoring for Stable Chronic Conditions

For patients with stable, well-controlled chronic conditions (for example, hypertension under pharmacotherapy or type 2 diabetes with HbA1c below 7.5%), longitudinal tracking of daily average vital parameters may provide adequate signal for trend detection. A clinician monitoring a stable COPD patient's resting heart rate trend over six months does not necessarily require 128 Hz data. A reliable daily average may suffice.



However, the transition from 'stable' to 'decompensating' is precisely when granularity becomes essential. A system built only for aggregate monitoring will miss the inflection point.

### 3.3 Synthesizing the Granularity Question: A Tiered Architecture

The resolution to the granularity debate is not binary. A well-designed Health OS should implement a tiered data architecture that matches data resolution to clinical use case.

Use Case Tier	Required Data Layer	Granularity Needed	Who Can Deliver
<b>Clinical alerting (AF, sepsis, deterioration)</b>	L0 - Raw PPG/ECG signal	Up to 128 Hz continuous, full waveform	<b>Hardware owners only (Corsano, Abbott)</b>
<b>Personalized predictive health</b>	L1 - Derived vital signs (consistent pipeline)	1 Hz minimum, continuous	<b>Full-stack hardware owners</b>
<b>Corporate wellness coaching</b>	L2 - Composite scores + L1 summaries	Daily / hourly summaries	Hardware owners + selected aggregators
<b>Population epidemiology</b>	L3 - Aggregated summaries	Daily averages sufficient	All platforms including HealthKit

Table 3. Tiered Data Architecture - Matching Resolution to Clinical Use Case

The critical insight from this framework is that a hardware sensor owner can operate at all four tiers simultaneously. It deploys different analytic pipelines against the same underlying raw data stream. An agnostic platform aggregator is structurally locked into Tiers 3 and 4. This creates an asymmetric capability envelope.

## 4. AI Training for a Health OS: Data Quality Over Data Volume

### 4.1 Why Large Language Models on Aggregated Health Data Fall Short

The emergence of large multimodal models capable of ingesting health data from HealthKit, Google Health and Samsung Health represents a genuine competitive development. Companies including Google (Med-PaLM, Gemini), Apple (health coaching features) and multiple well-funded startups are building LLM-based health insight engines on aggregated platform data. Understanding the specific limitations of this approach is essential for positioning the hardware-owner advantage.

LLMs trained on aggregated health platform data face four structural limitations that cannot be overcome through model scaling.

- **Data layer ceiling.** The highest-resolution physiological data these models can ingest is L3 (aggregated summaries) or at best L2 (composite scores from individual manufacturers). The L0 raw signals that contain the morphological disease signatures discussed in Section 3.1 are inaccessible.
- **Manufacturer confounding.** As shown in Table 2, 'HRV' means different things depending on device manufacturer. A model trained on multi-manufacturer HRV from HealthKit is learning device-specific algorithmic artifacts as much as true physiological variation.
- **Temporal resolution limitation.** HealthKit data granularity for most metrics is 1-minute intervals at best and 1-hour intervals at typical usage. Physiological events relevant to clinical outcomes occur at second-to-millisecond resolution. The model cannot detect what it cannot observe.



- Absence of ground truth linkage. Training a health AI to predict meaningful clinical outcomes (hospitalization, AF episode, sepsis onset) requires paired high-resolution sensor data and ground truth clinical outcome labels. This ground truth is available to medical device companies conducting clinical studies with consenting patients. It is not available to platform aggregators who receive anonymized, pre-processed data.

## 4.2 The Training Data Advantage of Full-Stack Hardware Ownership

A hardware sensor owner with an MDR/FDA-cleared clinical device occupies a fundamentally different position in the AI training data ecosystem. The combination of clinical-grade data collection and regulatory authorization to conduct clinical studies creates access to the most valuable training data in health AI.

- Labeled clinical outcome data. In a hospital RPM study conducted with a Corsano device, each patient's raw physiological time series is paired with clinical outcome data including discharge diagnoses, vital sign measurements by clinical staff, medication records and adverse events. This labeled dataset is the foundation for training models that predict real clinical outcomes rather than lifestyle wellness scores.
- Algorithmic consistency. All patients in a Corsano study contribute data processed through an identical, version-controlled algorithmic pipeline. There is no manufacturer confounding. A heart rate computed for patient A and patient B uses the same algorithm, the same artifact rejection logic and the same averaging window.
- Longitudinal depth. The continuous monitoring capability of the CardioWatch 287 (up to 1/128 second resolution, 24 hours per day) creates longitudinal data records of extraordinary depth. A seven-day hospital study generates approximately 77 million PPG samples per patient. This temporal depth enables detection of physiological patterns that are invisible in intermittent monitoring.
- Clinical validation pipeline. The Corsano MDR/FDA regulatory process requires clinical validation studies that simultaneously generate peer-reviewable evidence and high-quality labeled training data. Each regulatory submission is also a data asset creation event.

## 4.3 The Compounding Data Flywheel

The strategic value of the raw data asset is not static. It compounds. Each additional patient monitored by a Corsano device adds to a proprietary training dataset that grows more valuable as it grows larger. Because the data is algorithmically consistent (same hardware, same firmware version, same algorithms), the  $n=10,000$  patient dataset is not ten thousand independent observations. It is ten thousand points in a consistent, high-dimensional physiological feature space that can be used to train models of progressively higher predictive power.

This is the essence of the Health OS flywheel: hardware feeds raw data, raw data improves AI models, better models produce better insights, better insights create device stickiness, stickiness drives more hardware sales and more hardware generates more raw data. An agnostic platform aggregator participates in a version of this flywheel but at a fundamental data quality disadvantage. They are training on the shadows of physiological data rather than the physiological data itself.

### Compounding Advantage

Corsano's data flywheel operates at two levels simultaneously: (1) clinical-grade data from hospital and RPM studies that trains high-value diagnostic AI and (2) executive coaching data from CardioMood users that trains lifestyle optimization models. The full-stack architecture means both data streams are algorithmically consistent and can be combined for cross-domain model training.



## 5. Hardware Owners vs. Agnostic Platforms

### 5.1 The Agnostic Platform Aggregators: Apple HealthKit, Google Health Connect and Samsung Health

The three dominant platform aggregators have built extraordinary scale. Apple HealthKit has over 1 billion users with Apple Watch data. Google Health Connect connects data from hundreds of third-party apps and devices. Samsung Health processes data from hundreds of millions of Galaxy devices. However, scale at the aggregation layer does not translate to depth at the data layer.

Apple HealthKit operates as a permission-mediated data sharing system that collects processed outputs from connected devices. Apple has made significant investments in clinical research via the Apple Heart Study and the Apple Women's Health Study, using the Watch ECG functionality. But those studies are conducted with Apple's own L1 ECG data, not with the aggregated HealthKit data from third-party devices. The HealthKit API exposes approximately 180 data types, virtually all at L2/L3 levels. Apple has no access to the raw PPG signal from an Oura Ring or a Corsano watch connected to HealthKit.

Google Health Connect and Samsung Health follow a similar architecture. Third-party data arriving at L2/L3 only.

### 5.2 Specialist Hardware Competitors: Oura, Whoop and Abbott

Oura, Whoop and Abbott Libre represent full-stack hardware positions that are architecturally similar to Corsano in data ownership terms but differ significantly in regulatory classification, clinical data access and target use cases.

Dimension	Corsano	Oura	Whoop	Abbott Libre	Apple Watch Ultra
FDA/MDR Cleared	Yes (MDR IIa + FDA 510(k))	No (wellness only)	No (wellness only)	Yes (FDA Class II + EU MDR)	Partial (ECG, AFib)
Raw data access	Full PPG at up to 128 Hz	Limited (no public API)	No raw access	Glucose raw signal	ECG + PPG (proprietary)
Clinical use validated	Yes: hospital RPM, AF, SpO2, NEWS2	Research only (no Rx claim)	Research only (no Rx claim)	Yes: diabetes management	Limited to AFib + ECG
HCP portal / RPM	Full clinical portal + patient app	No HCP portal	No HCP portal	LibreView HCP portal	No dedicated RPM portal
Continuous monitoring	Up to 128 Hz (1/128 second)	Intermittent + sleep	Continuous HR (no PPG export)	Continuous glucose (15-min)	Intermittent (on-demand)
Full stack ownership	HW+FW+Algo+App+Cloud+Portal	HW+FW+Algo+App+Cloud	HW+FW+Algo+App+Cloud	HW+FW+Cloud+Portal	HW+FW+Algo+App+Cloud

Table 4. Competitive Comparison - Full-Stack Health OS Capabilities

The competitive analysis reveals that Corsano occupies a uniquely comprehensive position. It is the only company in the comparison that simultaneously offers FDA/MDR regulatory clearance across multiple vital signs, full raw PPG data access at up to 128 Hz, a clinical HCP portal for RPM workflows and full-stack ownership from hardware to cloud. Oura and Whoop, despite their consumer brand recognition and scale, are structurally excluded from clinical Health OS applications by their wellness-only regulatory classification. Abbott Libre holds both FDA and EU MDR clearance but is a single-parameter system (glucose) targeting a defined chronic disease category.



## 6. The Regulatory Moat: Clinical Grade vs. Consumer Wellness

---

### 6.1 The Clinical Grade Regulatory Framework as Competitive Barrier

Medical device regulation under the EU Medical Device Regulation (MDR 2017/745) and the U.S. FDA 510(k) framework represents the most durable competitive moat available to a Health OS platform company. Regulatory clearance is not merely a compliance exercise. It is a multi-year, multi-million euro investment in clinical evidence generation, quality management systems, post-market surveillance infrastructure and technical documentation that competitors cannot shortcut.

Corsano's CardioWatch 287-2 holds MDR Class IIa classification. This classification requires conformity assessment by a notified body (KIWA), clinical evaluation reports, post-market clinical follow-up plans and post-market surveillance reports (PSUR). The FDA 510(k) clearance requires demonstration of substantial equivalence to a predicate device across all intended use claims, supported by clinical validation data. These regulatory assets create a legal boundary around the clinical use case that consumer wellness companies cannot cross without equivalent investment.

The practical implication for Health OS positioning is decisive. An agnostic platform aggregator or a wellness-classified wearable company cannot make diagnostic or monitoring claims in regulated clinical settings regardless of the sophistication of their AI models. A Whoop or Oura device, however technically capable, cannot be prescribed for AF monitoring, cannot be reimbursed under RPM billing codes and cannot be deployed in an ICU workflow. Corsano can. It has the clinical evidence, quality systems and regulatory approvals to substantiate these claims in both European and U.S. markets.

### 6.2 Reimbursement as the Revenue Moat

The U.S. CMS Remote Patient Monitoring (RPM) billing framework (CPT codes 99453, 99454, 99457 and 99458) creates a reimbursement pathway that is exclusively available to FDA-cleared physiological monitoring devices. Under these codes, a physician practice or health system can bill Medicare and commercial insurers for the setup, data collection and clinical management services associated with continuous monitoring of patients with chronic conditions such as hypertension, AF, CHF, COPD and diabetes.

The RPM reimbursement value per patient per month (approximately \$100 to \$150 for a fully utilized monitoring program) represents a durable, recurring revenue stream that is unavailable to consumer wellness devices. A Health OS built on Corsano's clinical-grade device can be embedded in a physician's RPM billing workflow, generating revenue with every 16-day monitoring period. No HealthKit app, no Whoop integration and no Oura-based wellness platform can access this revenue stream. The regulatory clearance is, in effect, a revenue license.

### 6.3 The EU MDR Rule 11 Position: AI-Based Diagnostics

A forward-looking dimension of the regulatory moat concerns the EU MDR's Rule 11 classification rule. This rule applies to software intended to provide information used to make decisions with diagnosis or therapeutic purposes. Under Rule 11, AI-based diagnostic algorithms embedded in a medical device may be classified as Class IIa or higher. That is the same classification held by the CardioWatch 287-2 hardware.



This creates a complex but advantageous regulatory architecture for Corsano. The AI algorithms embedded in the CardioWatch 287-2 system carry regulatory status as part of the cleared device. A software-only health AI company attempting to make equivalent claims would need to qualify its algorithm as a Software as a Medical Device (SaMD). This process is equivalent in rigor to a full hardware device submission. Corsano's hardware regulatory clearance provides an umbrella for the AI inference layer that competitors building pure-software Health OS platforms must obtain independently.

## 6.4 EU AI Act Compliance: A Further Structural Advantage

The EU Artificial Intelligence Act (AI Act), which entered into force in August 2024, introduces binding requirements for AI systems deployed in high-risk contexts. Medical AI systems that influence clinical decision-making fall squarely into the high-risk category under Annex III of the Act. Compliance requires documented risk management, data governance frameworks, human oversight provisions, robustness testing and transparency obligations aligned with the overall quality management system.

Corsano has fully implemented the EU AI Act requirements into its Quality Management System (QMS). This means that every AI-driven alert, predictive model and automated decision support function within the CardioWatch 287-2 platform is governed by documented AI risk management procedures, training data governance policies and human oversight protocols that satisfy the Act's requirements. This is a significant compliance achievement: most competitor platforms, including wellness-classified wearables and agnostic aggregators deploying AI-based recommendations, have not yet achieved equivalent compliance status.

The practical implication is twofold. First, Corsano's AI-driven clinical features can be deployed in EU healthcare settings without incremental regulatory exposure, because the AI governance framework is already embedded in the QMS. Second, as the AI Act's enforcement provisions take effect through 2026 and 2027, competitors that have not implemented compliant AI governance frameworks will face either enforcement risk or costly remediation programs. Corsano's early full compliance converts a regulatory burden into a competitive moat.

### EU AI Act Compliance

Corsano is among the first wearable medical device companies to achieve full EU AI Act implementation within its QMS. This positions every AI-driven clinical feature of the CardioWatch 287-2 platform for immediate deployment in EU regulated healthcare settings, without incremental compliance overhead. Competitors face enforcement risk as EU AI Act enforcement escalates through 2027.



## 7. Corsano's Full-Stack Architecture: The Competitive Synthesis

### 7.1 The Seven Layers of the Corsano Health OS Stack

The Corsano Health OS value proposition rests on the integration of seven architectural layers, each of which reinforces the others and creates switching costs that protect the platform position. Unlike point solutions or aggregation platforms, Corsano's architecture creates value through the interactions between layers. Economists call this 'complementarity.' Strategists call it 'integration advantage.'

#	Layer	Corsano Component	Strategic Value
1	Hardware	CardioWatch 287-2: 3-wavelength PPG, 3-axis accelerometer, heatflux, impedance. MDR IIa / FDA 510(k) cleared	Owns the data source. Controls sensor geometry, wavelength selection and sampling architecture
2	Firmware	Proprietary embedded firmware: sampling control, on-device preprocessing, power management	Controls sampling rate (up to 128 Hz), artifact mitigation at source and data fidelity
3	Algorithms	Proprietary DSP algorithms: AF detection, SpO2, HR, HRV, RR, sleep staging, NEWS2 and Cor-Events alert framework (35 defined events)	Creates L1 vital parameters with algorithmically consistent methodology. Enables reprocessing of L0 with updated algorithms
4	Patient App	iOS/Android patient-facing application: data display, alert notifications, medication adherence and patient-reported outcomes	Generates behavioral and PRO data layer that enriches physiological signals. Creates patient engagement and adherence flywheel
5	Cloud Platform	Secure HIPAA/GDPR-compliant cloud infrastructure: raw signal storage, real-time processing pipeline, API layer and data governance	Stores L0 raw data. Enables retrospective reanalysis with improved algorithms. Provides the AI training data repository
6	HCP Portal	Web-based clinical portal: real-time patient dashboards, Cor-Events alerts and trend analytics	Creates clinical workflow integration, labeled clinical outcome data for AI training
7	Coaching Platform	CardioMood: B2B executive coaching dashboard for enterprise Leadership & Culture.	Expands data asset into non-clinical population.

Table 5. Corsano Health OS - Full-Stack Architecture and Strategic Value by Layer

### 7.2 The 128 Hz Continuous Data Advantage: Quantifying the Signal Richness

Corsano's CardioWatch 287-2 continuous monitoring at up to 128 Hz generates data that is qualitatively different from consumer wellness wearables. The following table makes this concrete.

Metric	Corsano 128 Hz	Typical Consumer Wearable (5-min avg)	HealthKit API (hourly avg)
PPG samples / minute	7,680	~1 (1 averaged value)	~0.017 (1 per hour)
PPG samples / day (24h)	11,059,200	~288 (5-min averages)	24
Beat-level RR interval resolution	~1ms (clinically relevant)	Not available	Not available
AF detection capability	Continuous + real-time alert	Intermittent / retrospective	Not possible
Pulse wave morphology analysis	Full waveform available	Not available	Not available
Data volume per 30-day study	~332M samples (compressed ~4GB)	~8,640 values (~50KB)	~720 values (<5KB)

Table 6. Data Volume and Capability Comparison - Corsano 128 Hz vs. Consumer Aggregated Data



## 8. Conclusion

---

The hypothesis that hardware wearable sensor owners hold structural advantages in building a Health OS with meaningful AI insights is validated, with important nuance. The advantage is real, significant and durable at the clinical tier, where access to raw physiological signals at high temporal resolution is a prerequisite for the AI capabilities that generate genuine clinical value: arrhythmia detection, early deterioration alerting and personalized predictive modeling.

The advantage is real but more modest at the wellness and population health tiers, where aggregate data is sufficient for the insight generation required. Agnostic platform aggregators including Apple HealthKit, Google Health Connect and Samsung Health can compete effectively at these tiers. Their scale advantages may overwhelm hardware-owner advantages in consumer wellness applications. This is not a concession that undermines the full-stack thesis. It is a calibration that directs investment toward the clinical tiers where the data moat is deepest.

Corsano Health B.V. occupies a uniquely strong position in this landscape for six compounding reasons. It generates raw data at up to 128 Hz from a clinically validated, MDR IIa and FDA 510(k)-cleared device. It controls the full algorithmic pipeline from L0 to L3. It owns the HCP portal that creates clinical workflow integration and generates labeled outcome data. It has a growing proprietary data asset that compounds with every deployed device. It has the Medtronic distribution partnership that accelerates clinical deployment at U.S. scale. And it is building the CardioMood wellness brand that extends data asset growth into the corporate wellness segment.

The agnostic platform aggregators, even those powered by large language models with access to millions of HealthKit records, are training on the shadows of physiological data. Corsano is training on the physiological data itself. In the long run of AI model development in health, this distinction is not marginal. It is the difference between a weather forecast and a weather report.

### Final Assessment

The data moat created by full-stack hardware ownership (raw PPG signal at up to 128 Hz through FDA/MDR-cleared algorithms, clinical portals and a growing labeled training dataset) is Corsano's most durable competitive asset. Protecting and compounding this moat is the strategic priority that all commercial, regulatory and product decisions should serve.

## Contact Corsano Health

**Website:** [www.corsano.com](http://www.corsano.com)

**Contact:** [corsano.com/contact-us](http://corsano.com/contact-us)



## References and Further Reading

---

### Clinical and Regulatory References

1. Perez MV, et al. (2019). Large-Scale Assessment of a Smartwatch to Identify Atrial Fibrillation. *New England Journal of Medicine*, 381(20), 1909-1917.
2. Bumgarner JM, et al. (2018). Smartwatch Algorithm for Automated Detection of Atrial Fibrillation. *Journal of the American College of Cardiology*, 71(21), 2381-2388.
3. National Institute for Health and Care Excellence (NICE). (2019). NEWS2 and deterioration: Identifying and responding to acute illness in adults in hospital. NICE Guidelines.
4. European Parliament. (2017). Regulation (EU) 2017/745 on Medical Devices (MDR). Official Journal of the European Union.
5. U.S. Food and Drug Administration. (2019). De Novo Request for Atrial Fibrillation Detection Software. FDA 510(k) Database.
6. U.S. Centers for Medicare and Medicaid Services (CMS). (2023). Remote Patient Monitoring Billing Guide: CPT Codes 99453, 99454, 99457 and 99458. CMS.gov.
7. Singhal A, et al. (2023). Wearable-based HRV: A systematic review of validity and methodology across consumer and clinical devices. *npj Digital Medicine*, 6, 149.
8. Shaffer F, Ginsberg JP. (2017). An Overview of Heart Rate Variability Metrics and Norms. *Frontiers in Public Health*, 5, 258.

### AI and Health Data Platform References

9. Singhal M, et al. (2023). Toward a Foundation Model for Clinical Time Series. arXiv preprint 2310.07204.
10. Tu T, et al. (2024). Towards Generalist Biomedical AI. *NEJM AI*, 1(3).
11. Goldberger AL, et al. (2000). PhysioBank, PhysioToolkit and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation*, 101(23), e215-e220.
12. Hannun AY, et al. (2019). Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. *Nature Medicine*, 25(1), 65-69.

### Market and Competitive Intelligence

13. Grand View Research. (2025). Wearable Medical Devices Market Size, Share and Trends Analysis Report.
14. IQVIA Institute. (2023). Digital Health Trends 2023: Innovation, Evidence, Regulation and Adoption. IQVIA Institute Report.
15. Deloitte Center for Health Solutions. (2024). Remote Patient Monitoring: From Pilot to Scale. Deloitte Insights.

